

367.40448X00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant(s): HANNUKSELA, Miska
Serial No.: Not yet assigned
Filed: August 9, 2001
Title: VIDEO CODING
Group: Not yet assigned

J1046 U.S. PTO
09/924502



LETTER CLAIMING RIGHT OF PRIORITY

Honorable Commissioner of
Patents and Trademarks
Washington, D.C. 20231

August 9, 2001

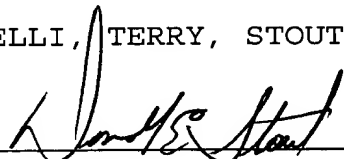
Sir:

Under the provisions of 35 USC 119 and 37 CFR 1.55, the applicant(s) hereby claim(s) the right of priority based on United Kingdom Patent Application No.(s) 0019983.6, filed August 14, 2001.

A certified copy of said United Kingdom Application is attached.

Respectfully submitted,

ANTONELLI, TERRY, STOUT & KRAUS, LLP



Donald E. Stout
Registration No. 26,422

DES/alb
Attachment
(703) 312-6600

THIS PAGE BLANK (USPTO)



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

j1046 U.S. PTO

09/924582



I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 19 July 2001

THIS PAGE BLANK (USPTO)

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)



The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

1. Your reference
PAT 00409 GB
2. Patent application number
(The Patent Office will fill in this part)
0019983.6 **14 AUG 2000**
3. Full name, address and postcode of the or of each applicant (underline all surnames)
NOKIA MOBILE PHONES LIMITED
KEILALAHDENTIE 4
02150 ESPOO
FINLAND

Patents ADP number (if you know it)

If the applicant is a corporate body, give the country/state of its incorporation
FINLAND

5911995004
4. Title of the invention
VIDEO CODING
5. Name of your agent (if you have one)
"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)
NOKIA IPR DEPARTMENT
NOKIA HOUSE
SUMMIT AVENUE
FARNBOROUGH
HAMPSHIRE
GU14ONG UK

Patents ADP number (if you know it) 7577638001 ✓
6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country	Priority application number (if you know it)	Date of filing (day / month / year)
7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application	Date of filing (day / month / year)
8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:
 - a) any applicant named in part 3 is not an inventor, or
 - b) there is an inventor who is not named as an applicant, or
 - c) any named applicant is a corporate body.
 See note (d))
YES

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description

22

Claim(s)

2

Abstract

1

Drawing(s)

7 x 1 m

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

1

Request for substantive examination (Patents Form 10/77)

Any other documents

(please specify)

11.

I/We request the grant of a patent on the basis of this application.

Signature

Juliet Hibbert Date 14/8/00

12. Name and daytime telephone number of person to contact in the United Kingdom

Juliet Hibbert Agent for the Applicant

JULIET HIBBERT 01252 865101

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

DUPLICATE

VIDEO CODING

This invention relates to video coding.

5

A video sequence consists of a series of still pictures or frames. Video compression methods are based on reducing the redundant and perceptually irrelevant parts of video sequences. The redundancy in video sequences can be categorised into spectral, spatial and temporal redundancy. Spectral
10 redundancy refers to the similarity between the different colour components of the same picture. Spatial redundancy results from the similarity between neighbouring pixels in a picture. Temporal redundancy exists because objects appearing in a previous image are also likely to appear in the current image. Compression can be achieved by taking advantage of this temporal
15 redundancy and predicting the current picture from another picture, termed an anchor or reference picture. Further compression is achieved by generating motion compensation data that describes the motion between the current picture and the previous picture.

20 However, sufficient compression cannot usually be achieved by only reducing the inherent redundancy of the sequence. Thus, video encoders also try to reduce the quality of those parts of the video sequence which are subjectively less important. In addition, the redundancy of the encoded bit-stream is reduced by means of efficient lossless coding of compression parameters and
25 coefficients. The main technique is to use variable length codes.

Video compression methods typically differentiate between pictures that utilise temporal redundancy reduction and those that do not. Compressed pictures that do not utilise temporal redundancy reduction methods are usually called
30 INTRA or I-frames or I-pictures. Temporally predicted images are usually

forwardly predicted from a picture occurring before the current picture and are called INTER or P-frames or P-pictures. In the INTER frame case, the predicted motion-compensated picture is rarely precise enough and therefore a spatially compressed prediction error frame is associated with each INTER
5 frame. INTER pictures may contain INTRA-coded areas.

Many video compression schemes also use temporally bi-directionally predicted frames, which are commonly referred to as B-pictures or B-frames. B-pictures are inserted between anchor picture pairs of I- and/or P-frames and
10 are predicted from either one or both of these anchor pictures. B-pictures normally yield increased compression compared with forward-predicted pictures. B-pictures are not used as anchor pictures, i.e., other pictures are not predicted from them. Therefore they can be discarded (intentionally or unintentionally) without impacting the picture quality of future pictures. Whilst
15 B-pictures may improve compression performance as compared with P-pictures, their generation requires greater computational complexity and memory usage, and they introduce additional delays. This may not be a problem for non-real time applications such as video streaming but may cause problems in real-time applications such as video-conferencing.

20 A compressed video clip typically consists of a sequence of pictures, which can be roughly categorised into temporally independent INTRA pictures and temporally differentially coded INTER pictures. Since the compression efficiency in INTRA pictures is normally lower than in INTER pictures, INTRA
25 pictures are used sparingly, especially in low bit-rate applications.

A video sequence may consist of a number of scenes or shots. The picture contents may be remarkably different from one scene to another, and therefore the first picture of a scene is typically INTRA-coded. There are
30 frequent scene changes in television and film material, whereas scene cuts

are relatively rare in video conferencing. In addition, INTRA pictures are typically inserted to stop temporal propagation of transmission errors in a reconstructed video signal and to provide random access points to a video bit-stream.

5

Compressed video is easily corrupted by transmission errors, for a variety of reasons. Firstly, due to utilisation of temporal predictive differential coding (INTER frames), an error is propagated both spatially and temporally. In practice this means that, once an error occurs, it is easily visible to the human eye for a relatively long time. Especially susceptible are transmissions at low bit-rates where there are only a few INTRA-coded frames, so temporal error propagation is not stopped for some time. Secondly, the use of variable length codes increases the susceptibility to errors. When a bit error alters a codeword, the decoder loses codeword synchronisation and also decodes subsequent error-free codewords (comprising several bits) incorrectly until the next synchronisation (or start) code. A synchronisation code is a bit pattern which cannot be generated from any legal combination of other codewords and such codes are added to the bit stream at intervals to enable re-synchronisation. In addition, errors occur when data is lost during transmission. For example, in video applications using the unreliable User Datagram Protocol (UDP) transport protocol in IP networks, network elements may discard parts of the encoded video bit-stream. Additionally, the size of an encoded INTRA picture is typically significantly larger than that of an encoded INTER picture. When periodic INTRA frames are encoded, it may therefore be more likely that an INTRA picture would be corrupted.

20
25

There are many ways for the receiver to address the corruption introduced in the transmission path. In general, on receipt of a signal, transmission errors are first detected and then corrected or concealed by the receiver. Error correction refers to the process of recovering the erroneous data perfectly as

30

if no errors had been introduced in the first place. Error concealment refers to the process of concealing the effects of transmission errors so that they are hardly visible in the reconstructed video sequence. Typically some amount of redundancy is added by the source or transport coding in order to facilitate error detection, correction and concealment.

In streaming applications, the demultiplexing-decompression-playback chain can be done while still downloading subsequent parts of the clip. In this case, a client may request retransmission of a corrupted INTRA picture. However, sometimes the streaming server may not be able to respond to such requests or the communication protocol between the streaming server and the client may not be able to transmit such requests. For example, the server may send a multicast stream, i.e., a single stream for multiple clients. If one of the clients receives a corrupted INTRA picture and sends a request for retransmission, the multicasting server would either send the retransmitted INTRA picture to all clients using the multicast channel or open an additional unicast channel for retransmission to the specific client. The former case increases network traffic unnecessarily for the majority of clients and the latter case complicates applications and network resource allocations. RTCP is an example of a transport protocol incapable of requesting retransmissions of specific pictures. Servers cannot determine which pictures to retransmit based on the receiver reports.

Current video coding standards define a syntax for a self-sufficient video bit-stream. The most popular standards at the time of writing are ITU-T Recommendation H.263, "Video coding for low bit rate communication", February 1998; ISO/IEC 14496-2, "Generic Coding of Audio-Visual Objects. Part 2: Visual", 1999 (known as MPEG-4); and ITU-T Recommendation H.262 (ISO/IEC 13818-2) (known as MPEG-2). These standards define a syntax for bit-streams and correspondingly for image sequences and images.

Some video coding schemes, such as ITU-T H.263 Annex N Reference Picture Selection (RPS) mode, include an indicator or indicator(s) to the prediction reference frame of an INTER frame.

5

Video Redundancy Coding (VRC) is also described in Annex N of H.263. In VRC, multiple "threads" of independently coded INTER-pictures are generated. Corruption of a frame in one thread will therefore not affect the decoding of the other threads. Periodically, all the threads converge to a so-called sync frame which may be an INTRA-picture or a non-INTRA picture that is redundantly represented within all threads. From this sync frame, all the independent threads are started again.

10

The principle of the VRC method is to divide the sequence of pictures into two or more threads in such a way that all pictures are assigned to one of the threads in a round-robin fashion. Each thread is coded independently. The frame rate within one thread is lower than the overall frame rate: half in the case of two threads, a third in the case of three threads and so on. This leads to a substantial coding penalty because of the generally larger changes and the longer motion vectors typically required to represent accurately the motion related changes between two P-pictures within a thread. At regular intervals all threads converge into a so-called sync frame. From this sync frame, a new series of threads is started.

15

20

If one of the threads is corrupted (e.g. because of packet loss), the remaining threads can be used to predict the next sync frame. It is possible to continue the decoding of the damaged thread, which typically leads to slight picture degradation, or to stop its decoding which leads to a decrease in the frame rate. If the length of the threads is kept reasonably small, however, both

25

forms of degradation will persist for a relatively short time, i.e. until the next sync frame is reached.

5 Sync frames are always predicted from one of the undamaged threads. This means that the number of transmitted I-pictures can be kept small, because there is no need for complete re-synchronization. Only if all threads are damaged between two sync frames, is correct sync frame prediction no longer possible. In this situation, annoying artefacts will be present until the next I-picture is decoded correctly, as would have been the case without employing
10 VRC.

According to a first aspect of the invention there is provided a method of encoding a video signal representing a sequence of pictures, the method comprising encoding at least a segment of a first picture of the sequence
15 without reference to another picture of the sequence and encoding at least the segment of said first picture with reference to another picture of the sequence to produce a corresponding temporally predicted picture or picture segment.

20 The invention is applicable whenever a frame or part of a frame is encoded in an INTRA manner. When the encoder encodes a frame (or part of a frame) in an INTRA manner, the encoder also encodes the frame (or part of the frame) again, this time in a temporally predictive manner with reference to another frame within the video sequence.

25 Preferably every picture (or picture segment) encoded without reference to another picture is also encoded with reference to another picture of the sequence.

The first picture (or picture segment) may be encoded with reference to another picture occurring in the sequence temporally prior to said first picture and/or with reference to another picture occurring in the sequence temporally after said first picture. The first picture (or picture segment) may be encoded
5 a plurality of times with reference to other pictures occurring in the sequence.

According to another aspect of the invention, there is provided a video encoder comprising an input for receiving a video signal representing a sequence of pictures, the encoder being arranged to encode a first picture (or
10 picture segment) of the sequence without reference to another picture of the sequence to produce a non-temporally predicted picture (or picture segment) and to encode said first picture (or picture segment) with reference to another picture of the sequence to produce a corresponding temporally predicted picture (or picture segment).

15 The invention also relates to a video codec including a video encoder according to the invention.

In a further aspect of the invention there is provided a multimedia system
20 including a video encoder according to the invention or a video codec according to the invention.

The invention also provides a video decoder and a method of decoding a video sequence in which a signal representing encoded pictures of a video
25 sequence is received, said video decoder being arranged to determine whether a non-temporally predicted frame has been corrupted and, if so, to monitor for a temporally-predicted representation of the frame and, on receipt of the temporally-predicted representation of the frame, to decode the temporally-predicted representation of the frame with reference to another
30 frame.

The invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

Figure 1 shows a multimedia mobile communications system;

5 Figure 2 shows an example of the multimedia components of a multimedia terminal;

Figure 3 shows an example of a video codec;

Figure 4 is a flow diagram illustrating the operation of a video encoder according to the invention;

10 Figure 5 shows an example of the output of a first embodiment of a video encoder according to the invention;

Figure 6 is a flow diagram illustrating the operation of a video decoder according to a first embodiment of the invention;

Figure 7 is a flow diagram illustrating the operation of a video encoder according to a second embodiment of the invention;

15 Figure 8 illustrates a further embodiment of the invention; and

Figure 9 shows a multimedia content creation and retrieval system into which a video encoder and/or decoder according to the invention may be incorporated.

20

Figure 1 shows a typical multimedia mobile communications system. A first multimedia mobile terminal 1 communicates with a second multimedia mobile terminal 2 via a radio link 3 to a mobile communications network 4. Control data is sent between the two terminals 1,2 as well as multimedia data.

25

Figure 2 shows the typical multimedia components of a terminal 1. The terminal comprises a video codec 10, an audio codec 20, a data protocol manager 30, a control manager 40, a multiplexer/demultiplexer 50 and a modem 60 (if required). The video codec 10 receives signals for coding from a video capture device of the terminal (not shown) (e.g. a camera) and receives

30

signals for decoding from a remote terminal 2 for display by the terminal 1 on a display 70. The audio codec 20 receives signals for coding from a microphone (not shown) of the terminal 1 and receive signals for decoding from a remote terminal 2 for reproduction by a loudspeaker (not shown) of the terminal 1. The terminal may be a portable radio communications device, such as a radio telephone.

The control manager 40 controls the operation of the video codec 10, the audio codec 20 and the data protocols manager 30. However, since the invention is concerned with the operation of the video codec 10, no further discussion of the audio codec 20 and protocol manager 30 will be provided.

Figure 3 shows an example of a video codec 10 according to the invention. The video codec comprises an encoder part 100 and a decoder part 200. The encoder part 100 comprises an input 101 for receiving a video signal from a camera or video source (not shown) of the terminal 1. A switch 102 switches the encoder between an INTRA-mode of coding and an INTER-mode. The encoder part 100 of the video codec 10 comprises a DCT transformer 103, a quantiser 104, an inverse quantiser 108, an inverse DCT transformer 109, an adder 110, one or more picture stores 107, a subtractor 106 for forming a prediction error, a switch 115 and an encoding control manager 105.

The decoder part 200 of the video codec 10 comprises an inverse quantiser 220, an inverse DCT transformer 221, a motion compensator 222, a plurality of picture stores 223 and a controller 224. The controller 224 receives video codec control signals demultiplexed from the encoded multimedia stream by the demultiplexer 50. In practice the controller 105 of the encoder and the controller 224 of the decoder may be the same processor.

The operation of an encoder according to the invention will now be described. The video codec 10 receives a video signal to be encoded. The encoder 100 of the video codec encodes the video signal by performing DCT transformation, quantisation and motion compensation. The encoded video data is then output to the multiplexer 50. The multiplexer 50 multiplexes the video data from the video codec 10 and control data from the control 40 (as well as other signals as appropriate) into a multimedia signal. The terminal 1 outputs this multimedia signal to the receiving terminal 2 via the modem 60 (if required).

10

In INTRA-mode, the video signal from the input 101 is transformed into DCT co-efficients by a DCT transformer 103. The DCT coefficients are then passed to the quantiser 104 where they are quantised. Both the switch 102 and the quantiser 104 are controlled by the encoding control manager 105 of the video codec, which may also receive feedback control from the receiving terminal 2 by means of the control manager 40. A decoded picture is then formed by passing the data output by the quantiser through the inverse quantiser 108 and applying an inverse DCT transform 109 to the inverse-quantised data. The resulting data is added to the picture store 107 by the adder 110, switch 115 being operated to present no data to the adder 110.

20

In INTER mode, the switch 102 is operated to accept from the subtractor 106 the difference between the signal from the input 101 and a reference picture stored in the picture store 107. The difference data output from the subtractor 106 represents the prediction error between the current picture and the reference picture stored in the picture store 107. A motion estimator 111 may generate motion compensation data from the data in the picture store 107 in a conventional manner.

25

The encoding control manager 105 decides whether to apply INTRA or INTER coding or whether to code the frame at all on the basis of either the output of the subtractor 106 or in response to feedback control data from a receiving decoder. The encoding control manager may decide not to code a received
5 frame at all when the similarity between the current frame and the reference frame is sufficient or there is not time to code the frame. The encoding control manager operates the switch 102 accordingly.

When not responding to feedback control data, the encoder typically encodes
10 a frame as an INTRA-frame either only at the start of coding (all other frames being P-frames), or at regular periods e.g. every 5s, or when the output of the subtractor exceeds a threshold i.e. when the current picture and that stored in the picture store 107 are judged to be too dissimilar to enable efficient temporal prediction. The encoder may also be programmed to encode frames
15 in a particular regular sequence e.g. I B B P B B P B B P B B P B B I B B P etc.

The video codec outputs the quantised DCT coefficients or prediction error data 112a, the quantising index 112b (i.e. the details of the quantiser used),
20 an INTRA/INTER flag 112c to indicate the mode of coding performed (I or P/B), a transmit flag 112d to indicate the number of the frame being coded and the motion vectors 112e for the INTER picture being coded. These are multiplexed together by the multiplexer 50 together with other multimedia signals.

25

Figure 4 is a flow diagram illustrating the operation of an encoder according to a first embodiment of the invention. The encoder creates temporally predicted frame(s) corresponding to INTRA coded frames as well as the INTRA coded frames. For simplicity the description only refers to the handling of INTER
30 and INTRA frames. The handling of other types of frames (e.g. B-frames) has

been omitted from the description. These are handled in a conventional manner.

5 Firstly the encoder receives a frame (401). The encoder decides (in a conventional manner) whether to code the frame or to skip it (402). If a decision is made to code the frame, the encoder decides (403) whether to code the frame as an INTER frame (404) or an INTRA frame (405). In the former case, the frame is coded (404) in a conventional temporally predictive manner. In the latter case, the frame is coded (405) in a conventional non-
10 temporally predictive manner. The decision to encode the frame as an INTRA frame (405) may be in response to a periodic INTRA request set up in the encoder or in response to a specific request received from a receiving decoder. Alternatively, it may be in response to the encoder determining that there is significant change in picture content which makes the use of an
15 INTRA picture beneficial from the point of view of compression efficiency. This may also be referred to informally as a scene cut.

The invention is applicable whenever a frame is encoded in an INTRA manner. When the encoder encodes a frame as an INTRA frame, the
20 encoder may also encode the frame again (407, 408), this time in a temporally predictive manner with reference to another frame within the video sequence.

For all but the first frame of the video sequence, a frame that has been INTRA-coded is also encoded (407) in a temporally predictive manner with
25 reference to a frame occurring within the video sequence temporally prior to the frame to be coded i.e. by forward prediction. This exception may also be the case for the first INTRA coded frame after a scene cut.

If the encoder determines (406) that the frame is the first INTRA frame to be
30 coded (or the first INTRA frame to be coded after a scene cut), the encoder

encodes the frame again (408), this time in a predictive manner with reference to a frame occurring within the video sequence temporally after the frame to be coded i.e. by backward prediction. The encoder therefore has to wait until a subsequent frame has been received for encoding before it can encode the predicted frame with reference to the subsequent frame.

In the example shown in Figure 4, the encoder waits (409) until the next INTRA frame has been encoded and then generates (408) the backward predicted representation of the earlier INTRA-coded frame with reference to the second INTRA frame of the scene.

All INTRA pictures after the first INTRA picture of a video sequence may have temporally predicted representation(s) encoded in a forward prediction manner and/or a backward prediction manner.

Figure 5 shows the frame-by-frame output of an encoder according to the invention. I0 is an INTRA coded picture appearing at the very beginning of the sequence (or after a scene cut). The system shown uses periodic INTRA pictures, and I4 is such a picture. In order to protect I0 and I4 from transmission errors, a second representation of each INTRA coded picture is temporally predictively encoded and transmitted. The second representation of picture 4 is an INTER picture predicted (and motion-compensated) with reference to picture 0. P4 could be predicted from any of the preceding frames, I0, P1, P2, or P3. Picture 0 is preferably chosen because otherwise P4 would be vulnerable to any transmission errors in the prediction path starting from I0. A receiving decoder can use any representation of picture 4 to reconstruct a picture for display/reference.

To protect picture 0 from transmission errors, a second representation of the same picture is coded and transmitted. Since picture 0 is the first INTRA

coded frame (or occurs immediately after a scene cut), the method of selection of a reference picture is different from that used for I4. Instead, the second representation of picture 0 is an INTER picture predicted (and motion-compensated) with reference to the picture that corresponds to the next periodic INTRA picture (picture 4 in the example). If I0 is corrupted, a receiver may wait until I4 and P0 have been received, after which it can reconstruct picture 0. Again P0 may be predicted from a picture other than frame I4, but prediction from the temporally nearest frame that has been coded in an INTRA manner is preferred.

10

The encoded data may be transmitted in the order in which it is encoded i.e. I0, P1, P2, P3, I4, P4, P0, P5... etc. Alternatively the frames of data may be re-ordered.

15 Considering the terminal 1 as receiving coded video data from terminal 2, the operation of the video codec 10 will now be described with reference to its decoding role. The terminal 1 receives a multimedia signal from the transmitting terminal 2. The demultiplexer 50 demultiplexes the multimedia signal and passes the video data to the video codec 10 and the control data to the control manager 40. The decoder 200 of the video codec decodes the encoded video data by inverse quantising, inverse DCT transforming and motion compensating the data. The controller 124 of the decoder checks the integrity of the received data and, if an error is detected, attempts to conceal the error in a manner to be described below. The decoded, corrected and
20 concealed video data is then stored in one of the picture stores 223 and
25 output for reproduction on a display 70 of the receiving terminal 1.

Errors in video data may occur at the picture level, the GOB level or the macroblock level. Error checking may be carried out at any or each of these
30 levels.

Figure 6 is a flow chart illustrating a video decoder according to the invention. Considering first the signal as illustrated in Figure 5, when a decoder according to the invention receives this signal each frame of the signal is
5 decoded in a conventional manner and then displayed on a display means. The decoded frame may be error corrected and error concealed in a conventional manner. Each time a frame is decoded, the decoder determines when the frame is to be displayed. This may be done by examining a temporal reference (TR) field of the header of the frame or, if the video is
10 transmitted in packets, the time-stamps of the packets may be examined.

Consider the frames shown in Figure 5, transmitted in the following order: I0, P1, P2, P3, I4, P4, P0, P5... etc. The decoder receives frame I0 (601) and determines (602) from its picture header that the frame is INTRA-coded. The
15 decoder decodes frame I0 (603) without reference to any other picture, displays it and stores it in picture store 223a (604). The decoder then receives frame P1 (601) and determines (602) from its picture header that the frame is INTER-coded as a P-frame. The decoder determines whether frame 1 has already been decoded (605). Since frame 1 has not already been
20 decoded, the decoder checks whether the reference frame for P1 has been received and decoded (606). As frame 0 has been received and decoded, the decoder decodes frame P1 (603) with reference to the reference frame 0, displays it and stores it in the next picture store 223b (604). The decoder then receives frame P2 (601) and determines (602) from its picture header
25 that the frame is INTER-coded as a P-frame. Frame 2 has not been decoded (605) and the reference frame 1 has been decoded (606) so the decoder therefore decodes frame P2 with reference to the preceding reference frame 1, displays it and stores it in the next picture store 223c and so on.

The next frame to be received is I4. The decoder then determines (602) from its picture header that the frame is INTRA-coded and decodes the frame. The decoder then receives frame P4 (601) and determines (602) from its picture header that the frame is INTER-coded as a P-frame. However the decoder
5 notes (605) that frame 4 has already been successfully decoded and discards frame P4 and receives the next frame of data P0. The decoder determines (602) from its picture header that the frame is INTER-coded as a P-frame. However the decoder notes (605) that frame 0 has already been decoded, discards frame P0 and subsequently receives the next frame of data P5. This
10 frame is then decoded with reference to decoded frame 4.

The above description describes what happens when no errors occur during transmission. Consider now the situation in which the decoder is unable to decode successfully the INTRA coded frame I0. This may be because the
15 data for the frame has been lost completely during transmission or some of the data for the frame has been corrupted in such a way that it is no longer recoverable satisfactorily. When frame P1 is received, the decoder determines that frame 1 has not yet been decoded (605) and that the reference frame for P1 has not been received and decoded (606). The
20 decoder then buffers the data for P1 (607). The same action is taken for P2 and P3.

When I4 is received successfully, the decoder decodes I4 without reference to any other frame. When P4 is received it is ignored because frame 4 has
25 already been decoded (605). When P0 is received the decoder determines that frame 0 has not already been decoded (605) and that the reference picture for P0 (frame 4) has been decoded (606). The decoder therefore decodes (603) frame P0 with reference to frame 4. Since frame 0 has now been decoded, it is possible to decode any buffered frames for which frame 0
30 was the reference frame. Thus the decoder decodes (608) the buffered frame

P1 with reference to frame 0. After decoding frame 1, the decoder is also able to decode the buffered frame P2 and, after decoding frame 2, the decoder is also able to decode the buffered frame P3.

- 5 The decoder discards the buffered frames if the decoder does not receive a reference frame for the buffered frames. The decoder may be arranged to discard the buffered frames if the reference frame for the buffered frames is not received within a given time or when an INTRA-coded frame is received.
- 10 Figure 7 is a flow chart also illustrating a second embodiment of a video decoder according to the invention, which is capable of detecting scene changes. The decoder receives (701) a signal for decoding. The decoder detects (702) if there has been a scene cut between the latest decoded frame and if the INTRA picture following the scene cut is missing. This may be done
- 15 for instance by examining a scene identifier in the header of the encoded frame or by examining the picture number of the frame, as set out in Annex W of H.263. In this latter case, the decoder can deduce from the picture number that a picture is missing and, if the next picture after a missed one does not indicate a spare reference picture number, the decoder assumes that the lost
- 20 picture is an INTRA picture associated with a scene cut.

If the INTRA picture immediately following a scene cut is not missing, the decoder checks (703) if the frame being decoded is a redundant representation of a frame that has already been decoded (called a sync frame

25 in the figure). In H.263, such a redundant representation may be indicated by giving the encoded frame the same TR as the primary representation of the picture.

If the frame is not a sync frame or if the primary representation of the sync

30 frame has not been decoded (704), the decoder decodes the frame (705) and

continues from the beginning of the loop. Otherwise, the decoder skips the decoding and continues from the beginning of the loop (501).

If the first INTRA picture (or that following a scene cut) is missing, the decoder starts to monitor (706) for the receipt of the backward-predicted frame corresponding to the first INTRA picture. The decoder knows that this particular frame should appear after the next periodic INTRA picture. All frames preceding the expected backward-predicted frame are buffered (707) (in compressed format). When the backward-predicted sync frame is due to appear in the bit-stream, the buffering is stopped. The next received frame is decoded (708) if it and the first periodic INTRA picture (i.e. the second INTRA picture in the scene) are received correctly. After that, all of the buffered frames are decoded (709), and then decoding can continue normally as if the first INTRA picture of the scene were never missing. Otherwise, the decoder cannot decode the buffered frames but rather it must discard them (710) and wait (711) for the next INTRA picture to arrive and continue decoding normally after decoding it.

Frame buffering has two implications. Firstly, the playback (displaying) process should use buffering to absorb the jitter introduced by the decoding process. This kind of a buffer is typical in video streaming and therefore the invention does not necessarily cause any modifications to the displaying process. Secondly, all buffered frames are decoded as soon as the reference frame is decoded. In practice this means that there may be a burst of computational activity when the first INTRA picture (I0) of a scene is lost and a subsequent representation (P0) of the picture of the scene is decoded. The decoder should be fast enough so that it has the resources to make up the time spent on buffering and so that uncompressed frames do not remain in the buffers for a long time.

The typical encoding/decoding order would be:

I0, P1, P2, ..., Pn-1, In, Pn, P0, Pn+1, Pn+2, ...

where I means an INTRA picture, P means an INTER picture, the index corresponds to relative capturing/displaying time, and time instant n corresponds to the first periodic INTRA frame after a scene cut.

The processing power requirement at the decoder can be minimised by re-ordering the transmitted frames compared with the typical encoding/decoding order, i.e., the transmitter can send the frames in the following order:

I0, In, P0, Pn, P1, P2, ..., Pn-1, Pn+1, Pn+2, ...

The invention is applicable to schemes where some segments of a picture are INTRA-coded and others INTER-coded. This is illustrated in Figure 8. Firstly, the bit-stream is coded in a conventional manner. When the encoder decides to update a picture segment in an INTRA manner, it carries out the following three operations:

1. The encoder codes the picture as a normal INTER-picture P(n).

2. The encoder codes P(n)', which is a redundant representation of P(n) where the desired picture segment is INTRA coded and other picture segments remain unchanged from P(n). The time-stamp of this newly coded picture is as close as possible to P(n).

3. The encoder encodes P(n)", which is a secondary representation of P(n)'. In P(n)", the picture segment of P(n)' that was coded in an INTRA manner is INTER-coded. Preferably the reference picture used is the latest picture (other than P(n)') in which the corresponding segment was INTRA-coded. The encoder may limit motion vectors to point only inside the picture segment in order to prevent the propagation of possible decoding errors from

the neighbouring segments. Alternatively, the encoder may use a coding mode such as H.263 Independent Segment Decoding (Annex R), which inherently prevents prediction from neighbouring picture segments. Other picture segments of $P(n)$ remain unchanged from $P(n)$.

5

After that, the encoder continues coding normally.

The decoder operates similarly to the one described earlier. If the picture segment in $P(n)$ is not corrupted, $P(n)'$ and $P(n)''$ are redundant and may be discarded. Otherwise, the decoder can use either $P(n)'$ or $P(n)''$ to recover the corrupted picture segment.

10

The error resilience of the invention can be improved if the transmitter sends multiple copies of a temporally predicted frame corresponding to an INTRA frame. To maximise error protection, these frames should be encapsulated into different packets. Moreover, it is beneficial if each INTRA picture is associated with both forward- and backward-predicted corresponding pictures.

15

The invention can be used in conversational, low-delay applications. In applications that do not buffer decoded frames before displaying them (but rather display decoded frames more or less immediately), corresponding backward predicted frames cannot be used by a decoder. Thus, a transmitting encoder for such applications (e.g. conversational applications) may be arranged not to code corresponding backward predicted frames for INTRA frames. Consequently, the first INTRA picture will not be protected in accordance with the invention.

20
25

The framework of an entire multimedia content creation and retrieval system will now be described with reference to Figure 9. The system has one or

30

more media sources 90, e.g. a camera and a microphone. Multimedia content can be synthetically created e.g. animated computer graphics and digitally generated music. To compose a multimedia clip consisting of different media types, the raw data captured from the sources are edited by an editing system 91. Typically the storage space required for raw multimedia data is huge. To facilitate an attractive multimedia retrieval service over low bit rate channels, multimedia clips are also compressed in the editing system 91. Then the clips are handed to a multimedia server 92. Typically, a number of clients 93 can access the server over a network. The server 92 is able to respond to requests presented by the clients 93. The main task for the server is to transmit a desired multimedia clip to a given client. The client 93 decompresses and plays the clip. In the playback phase, the client utilises one or more output devices 94, e.g. the screen and the loudspeaker of the client. In the system shown in Figure 9, The server incorporates a video encoder according to the invention and the client incorporates a video decoder according to the invention. In applications involving two-way video transmission, both the server and the clients incorporate a video codec according to the invention.

In a mobile multimedia retrieval system, at least part of the link connecting a client 93 and the server 92 is wireless e.g. by radio.

Typically, multimedia servers 92 have two modes of operation, namely they deliver either pre-stored multimedia clips or live (real-time) multimedia streams. In the former case, clips are stored on a server database, which is then accessed on an on-demand basis by the server for the client(s). In the latter case, multimedia clips are handed to the server as a continuous media stream that is immediately transmitted to clients 93. A server can remove and compress some of the header information produced by a multiplexing format as well as encapsulate the media clip into network packets.

The clients 93 control the operation of the server 92 using a control protocol. The minimum set of controls consists of a function to select a desired media clip. In addition, servers may support more advanced controls. For example, clients may be able to stop the transmission of a clip, to pause and resume the transmission of a clip, and to control the media flow if the throughput of the transmission channel varies, in which case the server dynamically adjusts the bit stream to fit into the available bandwidth.

- 10 The client 93 receives a compressed and multiplexed media clip from the multimedia server 92. First, the client demultiplexes the clip in order to get retrieve the separate media tracks and then decompresses these media tracks. After that, the decompressed (reconstructed) media tracks are played on output devices 94. In addition to these operations, a controller unit
- 15 interfaces with an end-user, controls the playback according to the user input and handles client-server control traffic. The demultiplexing-decompression-playback chain can be carried out while still downloading a subsequent part of the clip. This is commonly referred to as streaming. Alternatively, the whole clip is downloaded to the client 93 and then the client demultiplexes,
- 20 decompresses and plays it.

The invention is not intended to be limited to the video coding protocols discussed above: these are intended to be merely exemplary. The invention is applicable to any video coding protocol using temporal prediction. The

- 25 addition of additional INTER frames corresponding to INTRA frames as discussed above introduces error resilience into the encoded signal and allows a receiving decoder to select alternative decoding options if part of the received signal is corrupted.

CLAIMS

1. A method of encoding a video signal representing a sequence of pictures, the method comprising encoding a first picture of the sequence without reference to another picture of the sequence and encoding said first picture with reference to another picture of the sequence to produce a corresponding temporally predicted picture.
2. A method according to claim 1 wherein every picture encoded without reference to another picture is also encoded with reference to another picture of the sequence.
3. A method according to claim 1 wherein said first picture is encoded with reference to another picture occurring in the sequence temporally prior to said first picture.
4. A method according to claim 1 wherein said first picture is encoded with reference to another picture occurring in the sequence temporally after said first picture.
5. A method according to claim 1 wherein said first picture is encoded a plurality of times with reference to other pictures occurring in the sequence.
6. A video encoder comprising an input for receiving a video signal representing a sequence of pictures, the encoder being arranged to encode a first picture of the sequence without reference to another picture of the sequence to produce a non-temporally predicted picture and to encode said first picture with reference to another picture of the sequence to produce a corresponding temporally predicted picture.
7. A video codec including a video encoder according to claim 6.

8. A multimedia system including a video encoder according to claim 6 or a video codec according to claim 7.

5 9. A method of encoding a video signal representing a sequence of pictures, the method comprising encoding a segment of a first picture of the sequence without reference to another picture of the sequence and encoding
10 at least said segment of said first picture with reference to another picture of the sequence to produce a corresponding temporally predicted picture segment.

10. A method of video decoding comprising receiving a signal representing encoded pictures of a video sequence, determining whether a non-temporally predicted frame has been corrupted and, if so, monitoring for a temporally-
15 predicted representation of the frame and, on receipt of the temporally-predicted representation of the frame, decoding the frame with reference to another frame.

11. A video decoder comprising an input for receiving a signal representing
20 encoded pictures of a video sequence, said video decoder being arranged to determine whether a non-temporally predicted frame has been corrupted and, if so, to monitor for a temporally-predicted representation of the frame and, on receipt of the temporally-predicted representation of the frame, to decode the
25 temporally-predicted representation of the frame with reference to another frame.

12. A portable electronic device incorporating a video encoder or video decoder as claimed in any claim 6 or claim 11.

ABSTRACT

VIDEO CODING

- 5 A method of encoding a video signal representing a sequence of pictures, the method comprising encoding a first picture (or segment of a picture) of the sequence without reference to another picture of the sequence to produce a non-temporally predicted picture (I0) and encoding said first picture (or segment of a picture) with reference to another picture (I4) of the sequence to
- 10 produce a corresponding temporally predicted picture (P4) or segment of a picture.

Fig 5

THIS PAGE BLANK (USPTO,

1/7




Fig 1

THIS PAGE BLANK (USPTO)

2/7




Fig 2

THIS PAGE BLANK (USPTO)

3/7




Fig 3

THIS PAGE BLANK (USPTO)

4/7

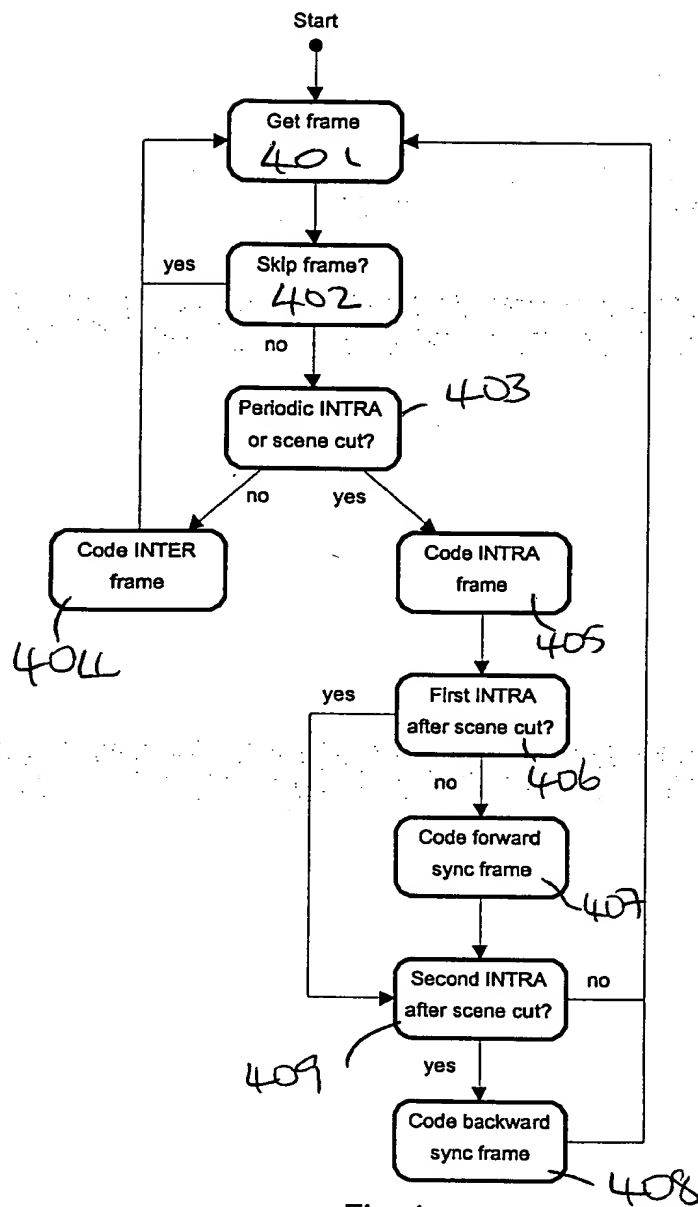


Fig. 4




Fig. 5

THIS PAGE BLANK (USPTO)

5/7

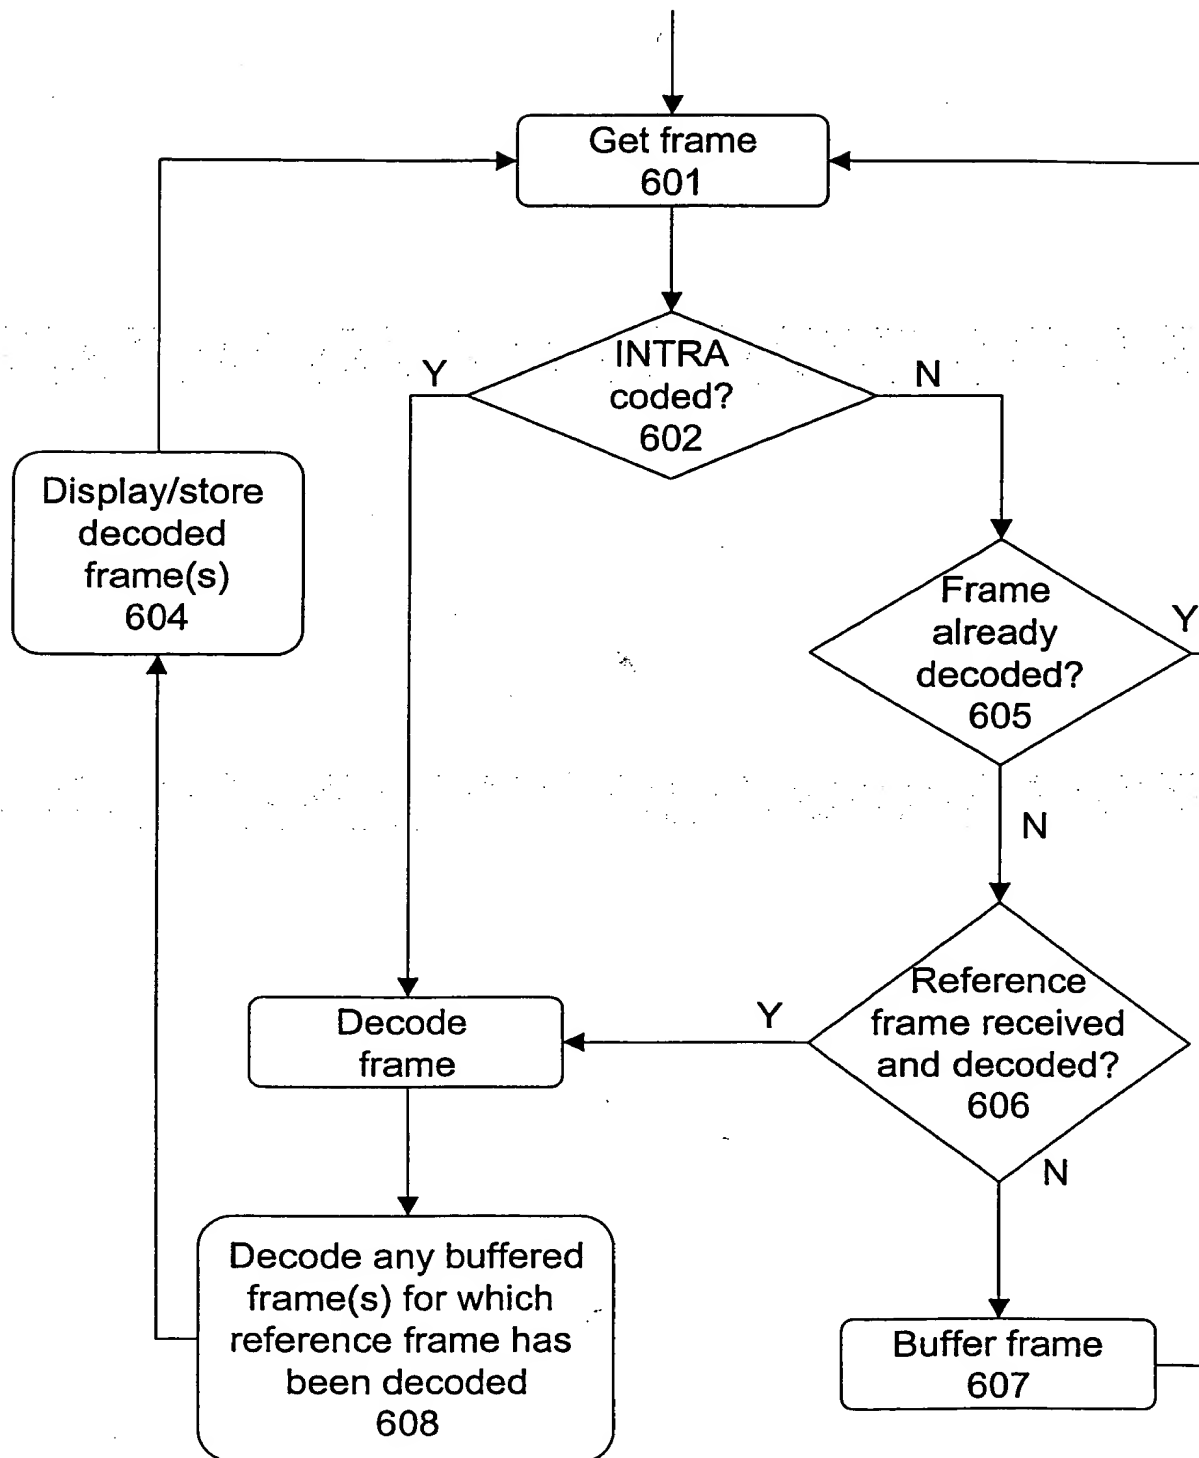


Fig. 6

THIS PAGE BLANK (USPTO)

6/7

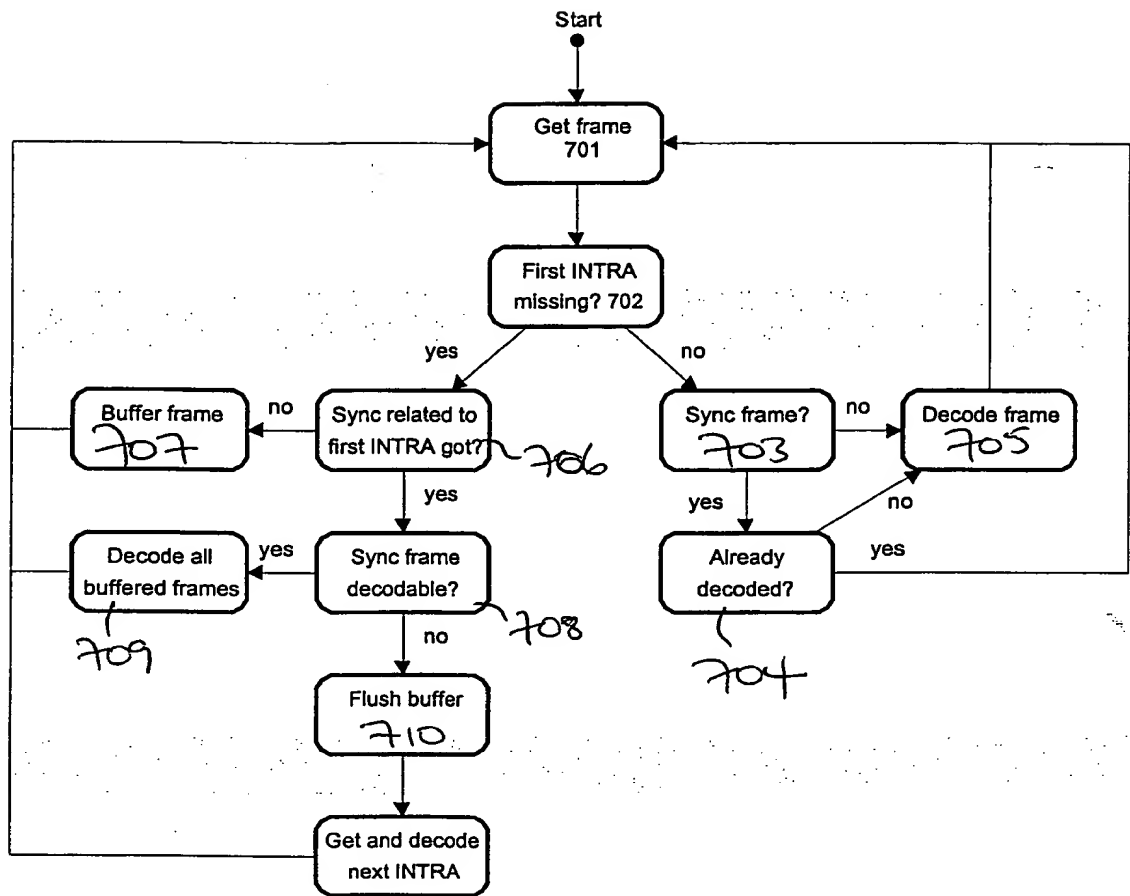


Fig. 7




Fig. 8

THIS PAGE BLANK (USPTO)

7/7

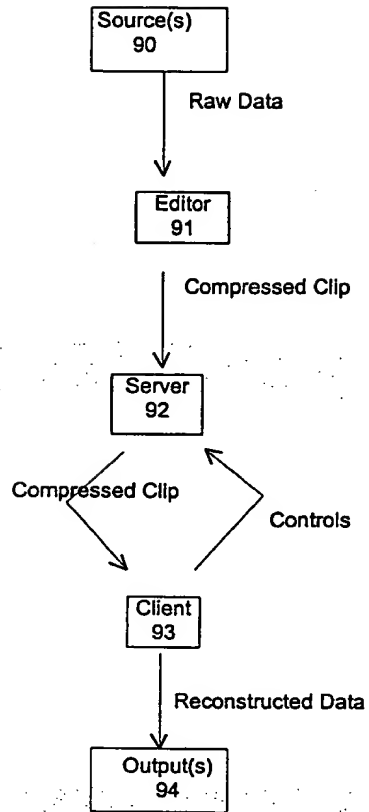


Fig. 9

THIS PAGE BLANK (USPTO)